

EV316936628

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Facilitating Performance Analysis for Processing

Inventor(s):
Kyle R. Johns
J. Andrew Goossen

ATTORNEY'S DOCKET NO. MS1-1671US

1 **TECHNICAL FIELD**

2 This invention relates to performance analysis, and more particularly to
3 facilitating performance analysis for processing.
4

5 **BACKGROUND**

6 As computing technology has advanced, the hardware on which video
7 games are run has also advanced and lead to increasingly complex video games.
8 Such complexity can result in video games with excellent graphics and a wide
9 range of actions that can be performed by the users. However, such complexity
10 can also make it difficult for game designers to create video games.

11 Problems or errors encountered in building and designing video games
12 oftentimes are resolved by "debugging" the instructions that make up the video
13 game. Unfortunately, such debugging can be a very tedious and time-consuming
14 process. For example, in order to maintain the desired playback rate for frames in
15 a video game (typically a rate of 30 frames per second or 60 frames per second),
16 each frame should require no more than a designated amount of time to be drawn
17 (e.g., at 60 frames per second, each frame should require no more than 16.67 ms
18 (milliseconds) to be drawn). If this designated amount of time is exceeded, then
19 the video playback can appear slow and the on-screen action may not appear
20 smooth. When the game designer sees portions of the video playback that are
21 slow, he or she typically tries to debug the instructions by finding out why the
22 designated amount of time is being exceeded, and changing the instructions so that
23 the certain amount of time is no longer exceeded.

24 Such debugging, however, can be very difficult. One reason for this
25 difficulty is that many video games are designed to be run on game systems that

1 include a graphics processing unit (GPU). The GPU receives commands to draw
2 various aspects of a scene, and renders the scene as indicated by these commands.
3 Unfortunately, the GPU frequently provides very little feedback, if any, to the
4 designer as to why the certain amount of time may be violated. As such, much of
5 the debugging effort is reduced to a trial and error methodology. Thus, it would be
6 beneficial to improve the quality and/or amount of information available to game
7 designers as they build and test their video games.

8 9 **SUMMARY**

10 Facilitating performance analysis for processing is described herein.

11 According to certain aspects, a set of commands to be submitted to a
12 processing unit is identified. A subset of the set of commands is selected and
13 submitted to the processing unit for processing. The processing performed by the
14 processing unit in response to the subset of the set of commands is analyzed.

15 According to other aspects, a stream of commands previously submitted to
16 a processing unit is identified. The stream of commands are modified and the
17 modified of commands are submitted to the processing unit. A difference between
18 a first amount of time required by the processing unit to process the stream of
19 commands and a second amount of time required by the processing unit to process
20 the modified stream of commands is determined.

21 According to other aspects, a state of a graphics processing unit is captured.
22 A plurality of commands submitted to the graphics processing unit in order to
23 draw a frame of video is also captured. Both the captured state and the captured
24 plurality of commands are saved.

BRIEF DESCRIPTION OF THE DRAWINGS

The same numbers are used throughout the document to reference like components and/or features.

Fig. 1 is a flowchart illustrating an example process for facilitating performance analysis for graphics processing.

Fig. 2 is a block diagram illustrating an example system for capturing commands sent to a graphics processing unit.

Fig. 3 is a block diagram illustrating another example system for capturing commands sent to a graphics processing unit.

Fig. 4 is a block diagram illustrating an example system for playback and analysis of a frame(s) drawn by a graphics processing unit.

Fig. 5 is a block diagram illustrating another example system for playback and analysis of a frame(s) drawn by a graphics processing unit.

Fig. 6 illustrates an example user interface that may be displayed.

Fig. 7 illustrates another example user interface that may be displayed.

Fig. 8 illustrates an example timeline window which could be displayed as a timeline window of Fig. 7.

Fig. 9 illustrates another example of an events window.

Fig. 10 illustrates an example of an images window that can be displayed as a frame window of Fig. 7.

Fig. 11 illustrates an example of the depth buffer view of an images window.

Fig. 12 illustrates an example of the wireframe view of an images window.

Fig. 13 illustrates an example of the texture view of an images window.

Fig. 14 illustrates an example of the overdraw view of an images window.

1 Fig. 15 illustrates an example of the fill rate view of an images window.

2 Fig. 16 illustrates an example of a call stack window that can be displayed
3 as a frame window of Fig. 7.

4 Fig. 17 illustrates an example of a warning window that can be displayed as
5 a frame window of Fig. 7.

6 Fig. 18 illustrates an example of a pushbuffer window that can be displayed
7 as a frame window of Fig. 7.

8 Fig. 19 illustrates an example of a summary window that can be displayed
9 as a frame window of Fig. 7.

10 Fig. 20 illustrates an example of a shaders window that can be displayed as
11 a frame window of Fig. 7.

12 Fig. 21 illustrates an example of a renderstates window that can be
13 displayed as a frame window of Fig. 7.

14 Fig. 22 illustrates an example of a texturestates window that can be
15 displayed as a frame window of Fig. 7.

16 Fig. 23 illustrates an example of an otherstates window that can be
17 displayed as a frame window of Fig. 7.

18 Fig. 24 illustrates an example pixel history window that can be displayed.

19 Fig. 25 illustrates an example pixel shader debugger window that can be
20 displayed.

21 Fig. 26 illustrates an example vertex shader debugger window that can be
22 displayed.

23 Fig. 27 illustrates an example mesh debugger window that can be
24 displayed.
25

1 Fig. 28 illustrates an example of a general computer environment, which
2 can be used to implement the techniques described herein.

3 4 **DETAILED DESCRIPTION**

5 Facilitating performance analysis for processing is described herein.
6 Although discussed herein primarily with reference to graphics processing and
7 graphics processing units, performance analysis for other types of processors can
8 also be facilitated. Commands submitted to a graphics processing unit (GPU) for
9 drawing a frame of video data are captured and saved. These captured commands
10 can subsequently be modified and submitted to the GPU, and/or only some of the
11 commands may be submitted to the GPU. By modifying the captured commands
12 appropriately, and/or submitting appropriate subsets of the captured commands, a
13 wide variety of information can be obtained regarding how the frame is being
14 drawn.

15 Fig. 1 is a flowchart illustrating an example process 100 for facilitating
16 performance analysis for graphics processing. Process 100 can be implemented in
17 hardware, software, firmware, or combinations thereof. Additionally, process 100
18 may be performed in a single device or alternatively in multiple devices.

19 In process 100, commands for drawing a frame of video data are captured
20 and saved. The video data is typically the video portion of a video game, such as
21 video games played on gaming consoles or computers. However, it is to be
22 appreciated that the video data may be associated with other types of applications
23 other than video games, such as educational or reference applications, simulation
24 or emulation applications, productivity or office applications, and so forth.
25

1 Video is typically made up of a series of individual images which are
2 played back sequentially. When played back at a fast enough rate, different
3 objects within these images appear to the human eye to move smoothly. Each one
4 of these individual images is referred to as a frame or a scene herein. Different
5 rates of playback can be used, and the rates of 30 frames per second or 60 frames
6 per second are commonly used with the NTSC (National TV Standards
7 Committee) standard.

8 Process 100 is initiated by a request, typically from a user, to initiate a
9 frame capture. In response to the request, the state of the graphics processing unit
10 is captured (act 102) when drawing of the current frame is finished. Given the
11 speed at which video frames are drawn, a request to capture a frame or scene is
12 typically received while another frame or scene is being drawn. Process 100 waits
13 until drawing of the current frame is finished so that all of the data desired to be
14 captured for drawing an entire frame can be captured.

15 The state of the graphics processing unit refers to all of the registers,
16 buffers, and/or other memory of the graphics processing unit in which variables or
17 settings are stored. These variables or settings are used by the graphics processing
18 unit in drawing frames. A variety of different variables or settings may be used,
19 such as shading or lighting variables, transparency variables, and so forth. Many
20 of these variables and/or settings can be programmed by the game designer, and
21 thus their settings are relevant to the frame about to be drawn. For example, a
22 frame being drawn may appear differently if transparency is turned on than if
23 transparency is turned off.

24 After the state of the graphics processing unit is captured, the commands
25 sent to the graphics processing unit for a frame are captured (act 104). The

1 number of commands sent to the graphics processing unit in order for the unit to
2 draw a single frame can vary, but can easily be in the hundreds if not thousands of
3 commands. Additional information associated with these commands (e.g., the
4 contents of memory referenced by the commands) may also be captured, as
5 discussed in more detail below. A set of commands sent to the graphics
6 processing unit is also referred to herein as a command stream or stream of
7 commands.

8 The captured state of the graphics processing unit from act 102 and the
9 captured set of commands from act 104 are saved (act 106). By saving the state of
10 the GPU and the captured set of commands, at a later time the GPU can be set to
11 the same state as when the set of commands were captured, and then selected ones
12 of the commands can be submitted to the GPU, as discussed in more detail below.

13 Alternatively, rather than capturing the commands or in addition to
14 capturing the commands, timing data may be captured in act 104 and saved in act
15 106. As discussed in more detail below, this timing data can identify, for example,
16 an amount of time taken to draw the frame and/or amounts of time taken to draw
17 different parts of the frame.

18 The set of captured commands can optionally be modified (act 108). This
19 modification can be changing one or more of the commands in the set, adding one
20 or more commands to the set, removing one or more commands from the set,
21 and/or reordering one or more commands in the set. The specific way in which
22 the modification is performed can vary based on the type of information to be
23 ascertained. For example, the modification may involve analyzing the set of
24 commands to identify redundant commands and then removing those redundant
25

1 commands. Additional details regarding modifications to the set of captured
2 commands are discussed below.

3 The GPU is then set to the state that was captured in act 102 (act 110). This
4 setting can be performed by, for example, setting all of the registers, buffers,
5 and/or other memory of the graphics processing unit in which variables or settings
6 are stored to the same settings as were captured in act 102.

7 Once the GPU is set to the captured state, at least a subset of the set of
8 commands captured in act 104 are sent to the GPU (act 112). These commands
9 sent to the GPU in act 112 may optionally include the set of commands as
10 modified in act 108. For example, if there are commands to draw the terrain in a
11 frame and commands to draw characters in the frame, only the commands to draw
12 the terrain may be sent to the GPU in act 112 to allow the user to see how the
13 terrain was drawn without the characters being present. By way of another
14 example, if the set of commands is modified in act 108 to remove the redundant
15 commands, then the modified set of commands (having the redundant commands
16 removed) may be sent to the GPU in act 112.

17 After the commands are sent to the GPU in act 112, feedback is returned
18 based on the commands sent in act 112 (act 114). This feedback can take a variety
19 of different forms. For example, the feedback may be the frame drawn given the
20 commands sent to the GPU in act 112. By way of another example, the feedback
21 may be an indication of how long it took the GPU to draw the frame given the
22 commands sent to the GPU in act 112. The feedback may also take other forms,
23 as discussed in additional detail below.

24 Fig. 2 is a block diagram illustrating an example system 150 for capturing
25 commands sent to a graphics processing unit. System 150 includes a game device

1 152 and a computing device 154. Game device 152 represents any of a wide
2 variety of game devices, such as stand-alone gaming consoles, handheld gaming
3 devices (e.g., with built-in display devices), cellular telephones, desktop PCs,
4 notebook or handheld PCs, and so forth.

5 Game device 152 includes a central processing unit (CPU) 156 and a
6 graphics processing unit (GPU) 158. CPU 156 runs a game or other application
7 160, and during running of the game 160 various commands are presented to GPU
8 158 in order to generate the video data for display by the game. Any of a variety
9 of conventional CPUs and GPUs can be used as CPU 156 and GPU 158,
10 respectively. GPU 158 processes the commands it receives and outputs its results
11 into a memory commonly referred to as a frame buffer. Once GPU 158 has
12 processed all of the commands it has received and output its results to the frame
13 buffer, the display device can display the frame using the contents of the frame
14 buffer. In certain embodiments, the frame buffer is a bit map that identifies, for
15 each pixel of the display device, what color is to be displayed for that pixel. The
16 processing of the commands and outputting of results into the frame buffer by
17 GPU 158 is also referred to as GPU 158 drawing the frame.

18 The video is displayed on one or more display devices 162, and the user
19 can control the game through one or more input devices 164. Display device 162
20 is intended to represent any device on which video can be displayed. Display
21 device 162 can be, for example, a television (e.g., cathode-ray tube (CRT),
22 plasma, projection, etc.), a computer monitor (e.g., CRT, plasma, LCD, etc.), a
23 display device built into the same housing as game device 152 (e.g., a plasma
24 display, LCD display, etc.), and so forth. Input device 164 is intended to represent
25 any device which can input data to the game device 152. Input device 164 can be,

1 for example, a handheld controller (e.g., having a joystick(s), button(s), and/or
2 triggers), a mouse or other cursor control device, a steering wheel and pedals, a
3 keyboard, a microphone, and so forth. Input device 164 can be external to game
4 device 152, or alternatively may be included as part of game device 152.

5 Computing device 154 is a remote device to game device 152. Computing
6 device 154 may be located close to game device 152 (e.g., on the same desk or in
7 the same room), or alternatively may be located further away from game device
8 152 (e.g., in another part of the building, or across the Internet). Computing
9 device 154 represents any of a variety of computing devices (e.g., a desktop PC,
10 workstation, portable or notebook computer, etc.). Computing device 154
11 similarly is coupled to (and/or includes) one or more display devices 166 and one
12 or more input devices 168. Display device 166 can be any device on which video
13 can be displayed, analogous to display device 162. Input device 168 can be any
14 device which can input data to computing device 154, analogous to input device
15 164.

16 Computing device 154 includes a capture control application 170 that is
17 executed by one or more processors (not shown) of computing device 154.
18 Capture control application 170 displays, on display device 166, a user interface
19 that allows a user (e.g., a game designer or tester) to request that frames of video
20 be captured and/or to present feedback (e.g., from act 114 of Fig. 1) to the user by
21 way of display device 166.

22 When a user, such as a game designer or tester, desires to capture a frame
23 of video data, he or she indicates to capture application 172 to begin capturing a
24 frame of video data. Capture application 172 can be invoked to begin capturing a
25 frame of video data in a variety of different manners. In certain embodiments, a

1 command is entered by the user to capture control application 170 by way of input
2 device 168 (e.g., a keyboard sequence may be entered, a "capture" button or menu
3 option displayed in the user interface provided by application 170 may be selected,
4 a verbal command may be input, and so forth). When such a command is received
5 by application 170, application 170 sends a request to capture application 172 for
6 capture application 172 to begin capturing a video frame. In other embodiments,
7 the command may be entered by the user to capture application 172 by way of
8 input device 164 (e.g., a keyboard sequence may be entered, a sequence of
9 buttons, triggers, and/or joystick positions on a game controller may be entered, a
10 verbal command may be input, and so forth).

11 The game designer or tester may desire to capture frames of video data at
12 different times for a variety of different reasons. For example, the designer or
13 tester may simply desire to randomly capture video frames while playing the game
14 in order to analyze the frames and verify they are being drawn as desired. By way
15 of another example, the designer or tester may play the game and, whenever he or
16 she perceives that the game is running slower than he or she believes it should,
17 then he or she can have a video frame selected at this "slow" spot. This allows, for
18 example, the designer or tester to analyze a frame in a "slow" spot of the game in
19 order to ascertain whether the frame is being drawn as fast as it should be, and if
20 not why not.

21 When capture application 172 receives the command to begin capturing a
22 video frame, application 172 waits for drawing of the current frame to finish, then
23 captures the state of GPU 158 and saves that state in capture storage 174. Capture
24 storage 174 can be any of a variety of storage devices (e.g., volatile or nonvolatile
25 memory such as RAM or Flash memory, a magnetic or optical disk, etc.).

1 Although illustrated as being part of game device 152, capture storage may be
2 located entirely or partially elsewhere. For example, capture storage 174 may be
3 distributed across multiple storage devices, or may be located entirely on a storage
4 device in another location. Other locations where capture storage 174 may be
5 located include an input device 164, computing device 154, input device 168, and
6 so forth.

7 After capturing the state of GPU 158, capture application 172 proceeds to
8 capture all commands that are sent to GPU 158 for the frame and stores the
9 captured commands in capture storage 174. Capturing the commands sent to GPU
10 158 is a process in which running application 160, which sends commands to GPU
11 158, is interrupted each time it sends a command and the command is copied into
12 capture storage 174. When the command has been copied, running application
13 160 continues to execute and submits its next command to GPU 158. After all of
14 the commands for the frame have been sent to GPU 158, capture application 172
15 stops capturing the commands. Capture application 172 then waits until it again
16 receives a command to capture a frame before proceeding to capture any
17 additional frames.

18 Although described herein as typically capturing only one frame at a time,
19 alternatively multiple frames may be captured by capture application 172 when it
20 receives a command to capture a frame. Capture application 172 may be pre-
21 configured or pre-programmed with a number of frames it will capture when it
22 receives a command to capture a frame, capture application 172 may have a
23 default number of frames it will capture, the user that initiates the command to
24 capture a frame may supply as part of the command the number of frames to be
25 captured, and so forth. For example, capture application 172 may capture two,

1 three, four, or more consecutive frames in response to such a command. By way
2 of another example, capture application 172 may capture one frame, then capture a
3 second frame some amount of time afterwards (e.g., the second frame may be the
4 45th frame after the command is received). By way of yet another example, a user
5 may specify when the capture is to begin and then frames are captured until the
6 user specifies that the capture should end. In this situation, a variable number of
7 frames would be captured.

8 Capture application 172 can capture timing data regarding commands
9 submitted to GPU 158 and/or the commands submitted to GPU 158. When timing
10 data and the commands are both captured, care should be taken that the time
11 involved in capturing of commands is not reflected in the captured timing data.
12 This situation can be resolved in different manners. In certain embodiments,
13 either the timing data or the commands are captured, and then the game 160 is
14 informed that it should repeat the frame. The game repeats preparing that frame at
15 which time the other of the timing data and commands are captured. In alternate
16 embodiments, the commands are captured and then timing data is obtained
17 subsequently as part of the analysis process (e.g., in acts 112 and 114 of Fig. 1).

18 The way in which the state of the GPU 158 and the commands sent to the
19 GPU 158 (or timing data) are captured can vary. The state of the GPU 158 and
20 commands sent to the GPU 158 may be captured in different manners due to
21 design choices as well as due to differences in the architecture of GPU 158 and/or
22 game device 152.

23 In certain embodiments, game device 152 includes a Direct3D® (D3D)
24 Application Programming Interface DLL (dynamic link library) which is a library
25

1 of graphics functions that can be invoked by game 160. Capture application 172
2 can optionally be incorporated into the D3D library.

3 Table I illustrates an example list of CPU events for which timing data can
4 be captured, as well as whether there is a GPU event that corresponds to the CPU
5 event. These events refer to calls to API (application programming interface)
6 functions or procedures. The time when the call to the API is made, as well as
7 when the called function or procedure ends, is recorded as part of the timing data.

Table I

CPU Event	API	GPU Event
Clear	Clear	Yes
DrawVertices	DrawVertices, DrawPrimitive	Yes
DrawVerticesUP	DrawVerticesUP, DrawPrimitiveUP	Yes
DrawIndexedVertices	DrawIndexedVertices, DrawIndexedPrimitive	Yes
DrawIndexedVerticesUP	DrawIndexedVerticesUP, DrawIndexedPrimitiveUP	Yes
CopyRects	CopyRects	Yes
BeginPush/EndPush	BeginPush/EndPush	Yes
RunPushBuffer	RunPushBuffer	Yes
Begin/End (Vertices)	Begin/End	Yes
Swap	Swap, Present	No
LockSurface	D3DSurface_LockRect	No
LockTexture	D3DTexture_LockRect	No
LockVertexBuffer	D3DVertexBuffer_Lock	No
LockPalette	D3DPalette_Lock	No
BlockOnObject	D3D block	No
BlockOnPushbuffer	D3D block	No
BlockOnFence	D3D block	No
BlockUntilIdle	D3D block	No
BlockOnSwap	D3D block	No
VBlank	None	No
VBlankSwap	None	No
VBlankMissed	None	No
PrimeVertexCache	PrimeVertexCache	Yes
DrawTri/RectPatch	DrawTri/RectPatch	Yes

Capture application 172 can obtain the current state of GPU 158 in a variety of manners. In certain embodiments, capture application 172 submits one or more commands specific to GPU 158, in response to which GPU 158 returns its state (e.g., the current values of all of the registers, buffers, and/or other memory of GPU 158 in which variables or settings are stored). Other static information in

1 GPU 158 may also be obtained as part of the capture, such as the instructions of
2 programs used internally by GPU 158 (e.g., pixel shader or vertex shader
3 programs used GPU 158). In other embodiments, the regions of memory in GPU
4 158 are accessible by memory address to capture application 172, and capture
5 application 172 can read from these memory addresses.

6 Capture application 172 can capture commands sent to GPU 158 by
7 monitoring the commands that are sent to GPU 158 (e.g., from the D3D DLL).
8 Every time a command is sent from the D3D DLL library, application 172
9 captures the command and all of the parameters of the command. Additionally,
10 some commands may have as parameters references to other memory locations.
11 In these situations, application 172 obtains the data from the referenced memory
12 location(s) and captures that data as well. Additionally, rather than always
13 capturing the data at the referenced memory location, application 172 can check
14 whether that memory location was previously referenced by another command
15 already captured for this frame – if so, then application 172 can check whether the
16 data previously obtained from that memory location is the same as the data
17 currently obtained from that memory location, and if so then the data need not be
18 captured again for this command. Such a check can be performed in different
19 manners, such as by maintaining a record of memory locations accessed and their
20 contents when accessed, by searching through the commands already captured,
21 and so forth.

22 By obtaining data from other memory locations referenced in calls to GPU
23 158, the capture performed by application 172 is more robust. As a specific
24 example, assume that texture mapping is being used to draw the surfaces of tree
25 leaves in a frame. The texture mapping for the leaves is stored in the same

1 memory location, but CPU 156 may modify that texture mapping for two different
2 leaves in the frame. So, even though each command sent to GPU 158 references
3 the same memory location, the texture mapping used for the two different leaves is
4 different. Thus, in order to accurately capture the frame being drawn, the data at
5 that memory location (the texture map) should be captured as well.

6 In certain embodiments, a capture opcode in the D3D DLL is used to
7 indicate whether the next frame is to be captured. When a command to capture a
8 frame is received by capture application 172, capture application 172 sets the
9 capture opcode. The end of a frame is indicated by game 160 by use of a Present()
10 or Swap() call. When a Present() or Swap() call is detected by the D3D DLL, the
11 D3D DLL checks whether the capture opcode has been set. If the capture opcode
12 has been set, then application 172 proceeds to begin capturing the next frame.
13 Additionally, application 172 clears the capture opcode so that at the end of the
14 frame being captured the capture opcode will not be set and another frame will not
15 be captured.

16 Additionally, the game designer may design the game to specify certain
17 information when the game is submitting commands to GPU 158. For example,
18 the game designer may include information identifying what part of the frame is
19 being drawn (e.g., the character's left arm, the character's right arm, the
20 character's weapon, the background terrain, etc.). When such information is
21 present, capture application 172 can capture this information and save it in capture
22 storage 174. This captured information can then subsequently be used when
23 feedback is presented to the user (e.g., in act 114 of Fig. 1). For example, a
24 timeline may be presented to the user in act 114 of Fig. 1 showing when different
25

1 parts of the frame were drawn, and the information identifying what part of the
2 frame is being drawn at any particular time can be shown on the timeline.

3 In addition, stack trace information may also be captured by application
4 172. Stack trace information tells the user which software routines called a
5 particular function (e.g., a D3D function) which resulted in a command being sent
6 to GPU 158. Each time a function is called, the return address is placed on the
7 stack. The stack trace is obtained by parsing back through the stack from the
8 current stack pointer and retrieving these return addresses. Using this series of
9 return addresses, capture application 172 examines the symbols for running
10 application 160 and converts the addresses to symbols that are more informative to
11 the user. For example, a fully resolved stack trace might look something like the
12 following:

```
13     Main()  
14     GameLoop()  
15     Render()  
16     DrawMainCharacter()  
17     FirstPass()  
18     RightArm()  
19     DrawIndexedVertices()
```

20 The last function listed in the stack trace, DrawIndexedVertices(), is
21 typically a function that resulted in a GPU command being passed to GPU 158.
22 When a developer is analyzing data from a captured scene, the stack trace helps
23 the developer to identify which drawing calls are which.

24 Fig. 3 is a block diagram illustrating another example system 200 for
25 capturing commands sent to a graphics processing unit. System 200 is similar to

1 system 150 of Fig. 1, except that system 200 includes a single game device and no
2 computing device.

3 Game device 200 includes a CPU 202 (analogous to CPU 156 of Fig. 2), a
4 GPU 204 (analogous to GPU 158 of Fig. 2), capture storage 206 (analogous to
5 capture storage 174 of Fig. 2), and game 208 (analogous to game 160 of Fig. 2).
6 Game device 200 is coupled to (or includes) display device 210 (analogous to
7 display device 162 or 166 of Fig. 2), and is also coupled to (or includes) input
8 device 212 (analogous to input device 164 or 168 of Fig. 2).

9 Game device 200 also includes a capture control application 214 and a
10 capture application 216. Capture control application 214 is similar to capture
11 control application 170 of Fig. 2, and Capture application 216 is similar to capture
12 application 172 of Fig. 2. However, applications 214 and 216 different from
13 applications 170 and 172 in that applications 214 and 216 are part of the same
14 game device 200. So, any commands or data sent between application 214 and
15 capture application 216 are sent internal to game device 200 rather than between
16 two devices. Additionally, capture control application 214 presents its user
17 interface by way of display device 210.

18 Fig. 4 is a block diagram illustrating an example system 240 for playback
19 and analysis of a frame(s) drawn by a graphics processing unit. System 240
20 includes a game device 242 and a computing device 244. Game device 242
21 includes a CPU 246 (analogous to CPU 156 of Fig. 2) and a GPU 248 (analogous
22 to GPU 158 of Fig. 2). Game device 242 is coupled to (or includes) display device
23 250 (analogous to display device 162 of Fig. 2), and is also coupled to (or
24 includes) input device 252 (analogous to input device 164 of Fig. 2). Similarly,
25 computing device 244 is coupled to (or includes) display device 254 (analogous to

1 display device 166 of Fig. 2), and is also coupled to (or includes) input device 256
2 (analogous to input device 168 of Fig. 2).

3 Game device 242 is similar to game device 152 of Fig. 2, and in certain
4 embodiments devices 242 and 152 are the same device (or at least the same type
5 of device, such as both the same type of gaming console). The difference between
6 game devices 242 and 152 is that game device 242 is running an application(s) to
7 facilitate playback and analysis of a frame based on previously captured
8 commands and GPU state (e.g., captured by capture application 172 of Fig. 2).

9 Computing device 244 includes an analysis control and feedback
10 application 258. Application 258 provides a user interface, through display device
11 254. The user interface allows the user of device 244 to input, through input
12 device 256, requests for analysis of a frame to be performed. These requests can
13 also be referred to as questions regarding the frame. The user interface provided
14 by application 258 also displays feedback regarding the frame to the user of device
15 244 (e.g., the feedback in act 114 of Fig. 1). The specific feedback provided to the
16 user can vary based on the analysis performed, as discussed in more detail below.

17 Requests for analysis are sent from application 258 to analysis application
18 260. Analysis of a frame can begin immediately after the frame is captured, or
19 alternatively at some later time. For example, when the designer or tester captures
20 a frame, the analysis of the frame may begin right after the frame is captured, or
21 alternatively multiple frames may be captured and the designer or tester (or some
22 other user) may begin analysis of the captured frames at some later time.

23 Analysis application 260 receives the requests for analysis and sends the
24 appropriate portions of captured data 262 to GPU 248. Feedback regarding the
25 drawing of the data sent to GPU 248 is then returned by analysis application 260

1 to application 258 for presentation to the user on display device 254. Depending
2 on the analysis requested, analysis application 260 may generate modified data
3 264 that is sent to GPU 248 in place of, or alternatively in addition to, captured
4 data 262.

5 The captured data 262 is the same data for the frame as was captured and
6 stored in the capture storage (e.g., storage 174 of Fig. 2). In certain embodiments,
7 the captured data is maintained in the game device and referred to when requests
8 are sent by application 258. In other embodiments, when analysis of a frame is to
9 begin, the captured data is sent to game device 242 by application 258. In other
10 embodiments, the portions of the captured data (and/or modified data) necessary
11 for the analysis are sent by application 258 along with the request.

12 Additionally, although analysis application 260 is shown as being part of
13 game device 242, portions of analysis application 260 may be implemented in
14 computing device 244. For example, any required modifications to the data may
15 be performed by computing device 244 rather than game device 242.
16 Additionally, some analysis can be performed by computing device 244 without
17 any help from game device 242. For example, the size of the captured data and
18 the number of commands in the data can be determined by computing device 244
19 by parsing through the command stream.

20 When analysis of a frame is to begin, the state of GPU 248 is set to the
21 previously captured state (e.g., as captured in act 102 of Fig. 1). The various
22 memory location and register values that were captured as part of the state of the
23 GPU (e.g., GPU 158 of Fig. 2) are written back to GPU 248. By writing these
24 values back to the GPU, the GPU will be in the same state as it was at the
25 beginning of the captured frame.

1 It should be noted that, during the analysis process, certain other memory
2 locations in game device 242 may be accessed. For example, the memory
3 location(s) for texture map(s) used in drawing the frame may be accessed. In such
4 situations where other memory locations are accessed, analysis application 260
5 writes the appropriate data (as captured) to those memory locations prior to their
6 being accessed by the captured commands.

7 One type of analysis that can be performed by analysis application 260 is
8 timing analysis. Timing analysis refers to determining an amount of time a
9 particular frame or portion of a frame takes to be drawn. Game device 242 has a
10 clock (not shown) that can be used to determine an amount of time taken for a
11 particular frame to be drawn. The commands for drawing the frame are sent to
12 GPU 248 by analysis application 260, so application 260 can identify the current
13 time of the clock when the first command is sent to GPU 248 (e.g., immediately
14 prior to sending the first command to GPU 248). Application 260 also sends a
15 command to GPU 248 to indicate that it has completed sending commands to GPU
16 248 (e.g., the Present() or Swap() commands discussed above). When GPU 248
17 has completed drawing the last command, GPU 248 returns an indication to the
18 application sending the commands to GPU 248 (analysis application 260 in this
19 instance) that the last command is completed. Application 260 can then identify
20 the current time of the clock. By knowing the time when drawing of the frame
21 began and the time when drawing of the frame was completed, the difference
22 between these two times can be readily determined in order to identify the amount
23 of time taken to draw the frame.

24 The amount of time taken to draw a portion of a frame can be determined in
25 an analogous manner. By accessing the clock, the time when drawing of the

1 portion begins and drawing of the portion is completed can be readily determined.
2 These different portions can correspond to, for example, the events identified
3 above in Table I.

4 The timing analysis can be used as a basis to provide various feedback to
5 the user. The amount of time taken for a frame or portion of a frame to be drawn
6 can be returned to the user. Additionally, other information can be ascertained
7 based on the timing analysis and this other information can be returned to the user.
8 This other information can be ascertained by analysis application 260 and/or
9 analysis control and feedback application 258.

10 A wide variety of information can be ascertained based on this timing
11 analysis. For example, which portions (e.g., objects) of the frame took the longest
12 or shortest time to draw can be ascertained.

13 The timing analysis can also be performed based on modified data 264.
14 This allows analysis application 260 to modify captured data 262 and then submit
15 the modified data to GPU 248. The time taken to draw both the unmodified data
16 and the modified data can be determined, and these two times used to determine
17 how much the modification affected the time taken to draw the frame. For
18 example, analysis application 260 may analyze captured data 262 to identify
19 redundant commands. Redundant commands within the set of commands that is
20 captured data 262 refers to duplicate commands that are present in the set of
21 commands but one or more of these duplicate commands can be removed without
22 altering the frame drawn by the set of commands. E.g., multiple commands to
23 turn off transparency may be present in the set of commands, even though no other
24 command in the set of commands operates to turn on transparency. Thus, all of
25 the commands to turn off transparency after the first command to turn off

1 transparency could be removed from the set of commands (the first command may
2 also be removed if there is no previous command or state of the GPU 248 that
3 would have caused transparency to be turned on).

4 It should be noted that some timing data for performing the timing analysis
5 may be captured when the frame is captured, rather than being determined during
6 the analysis process. For example, situations may arise where the amount of time
7 taken by the CPU in the game device to determine particular values to be used in
8 drawing the frame is very large. By capturing the timing data when the CPU is
9 determining the particular values, rather than capturing the timing data based on
10 the resultant captured values, situations where the GPU goes idle can be more
11 easily detected. For example, the CPU can be so busy that it waits too long to
12 issue a new command to the GPU which causes the GPU to go idle because it has
13 nothing to work on. This is an important case for developers to be aware of
14 because they typically endeavor to keep the GPU busy at all times. This
15 information is lost when playing back a captured stream of commands because the
16 commands are played back one-after-another without regard for any idle spots that
17 may have been present when the application (e.g., game) originally ran.

18 Another type of analysis that can be performed by analysis application 260
19 is a partial frame draw. All of the commands in the captured data 262 are used to
20 draw the entire frame. A partial frame draw can be performed by sending only a
21 subset of the commands in the captured data 262 to GPU 248 and having GPU 248
22 draw the frame using only that subset (and thus less than all of the captured
23 commands). For example, analysis application 260 may choose to send only the
24 first five commands or the first ten commands, or all of the commands leading up
25 to a particular command, etc. to GPU 248.

1 Whatever the subset of commands is, GPU 248 receives the subset and
2 draws a frame given only that subset. Analysis application 260 can then return
3 this frame to analysis control and feedback application 258 for display to the user
4 on display device 254. Application 260 can send this frame in different manners,
5 such as generating a bitmap representing the frame as drawn by GPU 248 and
6 sending that bitmap to application 258. Drawing a partial frame can be useful to
7 the developer as it allows the developer to see what the frame looks like after the
8 first n drawing commands are completed (where the value of n can be set by the
9 developer). Much like single-stepping a program, this allows the developer to see
10 the scene being drawn step-by-step and to quickly determine which step drew
11 something incorrectly.

12 Another type of analysis that can be performed by analysis application 260
13 is the determining of internal non-exposed state of GPU 248. Some state of the
14 GPU 248 can be accessed by devices external to GPU 248 (e.g., the settings of
15 certain variables, register values, etc.), and these states are captured as discussed
16 above (e.g., in act 102 of Fig. 1). These states are typically states that affect all
17 future drawing commands on GPU 248.

18 However, GPU 248 may also include additional internal state for which
19 GPU 248 provides no mechanism for any component or module external to GPU
20 248 to access (e.g., no command can be sent to GPU 248 to read particular internal
21 register values). Such states are typically transitory states. Transitory states can
22 be thought of as similar to scratchpad memory, where the contents are only valid
23 during a particular instruction of a particular invocation of the vertex or pixel
24 shader program. When the next vertex or pixel is processed, this state is
25 overwritten. Such transitory states are typically not easily retrievable because

1 GPUs typically have no way to halt execution at a particular point so that these
2 states can be read.

3 Some GPUs include internal programs that process some of the data
4 received as commands. For example, one or more vertex shader programs and/or
5 one or more pixel shader programs may be included in the GPU that operate on
6 some of the data received by the GPU. As these programs within the GPU can
7 make use of their own internal variables (e.g., registers and/or other memories), it
8 may be useful to the designer or tester to know what the values of these registers
9 and/or variables are as each pixel or vertex is processed. These registers and/or
10 variables are often transitory states, such as the contents of a temporary register
11 after a particular instruction has executed in a pixel shader program of GPU 248
12 that is a particular pixel in a particular drawing operation in the scene (e.g., the
13 tmp0 register contents after instruction 3 has executed in the pixel shader program
14 that is processing pixel 238,242 in the 7th drawing operation in the scene). By
15 modifying the captured commands, and possibly internal programs for the GPU,
16 additional information can be obtained.

17 Such modifications to identify settings of internal registers and/or other
18 variables is illustrated by the following example. Assume that a pixel shader
19 program used by the GPU is as follows (where t0, t1, and r0 are internal GPU
20 registers):
21
22
23
24
25

```

1 tex t0 // load the result of texture stage t0 into register t0
2 tex t1 // load the result of texture stage t1 into register t1
3 add r0, t0, c0 // r0 = t0 + c0
4 mul r0, r0, t1 // r0 = r0 * t1
5 // r0 is the result written to the frame buffer

```

Analysis application 260 can determine the internal value of the register r0 in the third instruction as follows. The shader programmer is modified to be as follows:

```

7 tex t0 // load the result of texture stage t0 into register t0
8 tex t1 // load the result of texture stage t1 into register t1
9 add r0, t0, c0 // r0 = t0 + c0
10 // r0 is the result written to the frame buffer

```

Notice that the last instruction which loaded register r0 with the value $r0 * t1$ was removed. At the end of the shader program, r0 will contain the value of $t0 + c0$ and that will be the result value that will be written to the frame buffer for that pixel. The modified command stream is fed to the GPU and the value for the pixel being tested is read from the frame buffer and can be returned to application 258 for presentation to the user. This value is the value of the r0 register in the third instruction.

The program can be modified again as follows:

```

19 tex t0 // load the result of texture stage t0 into register t0
20 tex t1 // load the result of texture stage t1 into register t1
21 mov r0, t1 // r0 = t1
22 // r0 is the result written to the frame buffer

```

Now, the value written by the GPU to the frame buffer will be the value of the t1 register in the second instruction. This process can be followed repeatedly to extract the values of all the internal registers.

1 In certain embodiments, analysis control and feedback application 258 has
2 multiple default requests or questions that it submits to analysis application 260.
3 These default requests or questions are pre-programmed into application 258.
4 Additionally, the user of computing device 244 may also request information that
5 involves one or more additional requests or questions be sent to analysis
6 application 260.

7 In certain embodiments, analysis application 260 and/or analysis control
8 and feedback application 258 analyze the requests sent to analysis application 260
9 and generate warnings for the user regarding rules or recommendations that were
10 violated by the frame. For example, each different type of GPU typically has its
11 own recommended programming practices. Tests to determine whether these
12 programming practices were violated can be programmed in to application 258 or
13 260, so that each time a frame is analyzed these tests can be performed to
14 determine whether the recommendations for programming the GPU were violated.
15 A variety of such recommendations may exist, such as not setting certain register
16 values if certain GPU functionality is not being used in a frame, issuing certain
17 commands in certain orders, having the frame buffer and depth buffer both in tiled
18 memory, not clearing the screen multiple times (e.g., clearing the screen after it
19 has already been cleared), not submitting redundant commands, and so forth.
20 Additionally, warnings can be given different priorities. For example, the
21 warnings can be ranked according to how much time adhering to their associated
22 recommendations would save for this frame. By way of another example,
23 particular warnings may be presented only if they satisfy certain constraints (e.g.,
24 display a warning to the user only if not violating the recommendation would have
25 allowed the frame to be drawn a threshold amount faster).

1 Analysis of the frame can continue until analysis application 260 is
2 informed to stop analyzing the frame. The user, for example, can select a "stop"
3 or "cancel" option on the user interface presented by application 258 to stop
4 analyzing the frame.

5 It should be noted that all of the analysis of the set of commands, whether
6 captured commands or modified commands are sent to GPU 248, is based on the
7 commands being performed by the same GPU as will be running the game. As the
8 actual GPU can be used, an emulator or simulator need not be used. By using the
9 actual GPU, the accuracy of the analysis is improved due to any approximations or
10 errors that may be introduced by an emulator or simulator. Additionally, any
11 timing measurements obtained by an emulator or simulator would have little, if
12 any, meaning.

13 Fig. 5 is a block diagram illustrating another example system 300 for
14 playback and analysis of a frame(s) drawn by a graphics processing unit. System
15 300 is similar to system 240 of Fig. 4, except that system 300 includes a single
16 game device and no computing device.

17 Game device 300 includes a CPU 302 (analogous to CPU 246 of Fig. 4), a
18 GPU 304 (analogous to GPU 248 of Fig. 4), captured data 306 (analogous to
19 captured data 262 of Fig. 4), and modified data 308 (analogous to modified data
20 264 of Fig. 4). Game device 300 is coupled to (or includes) display device 310
21 (analogous to display device 250 or 254 of Fig. 4), and is also coupled to (or
22 includes) input device 312 (analogous to input device 252 or 256 of Fig. 4).

23 Game device 300 also includes an analysis control and feedback application
24 314 and an analysis application 316. Analysis control and feedback application
25 314 is similar to analysis control and feedback application 258 of Fig. 4, and

1 analysis application 316 is similar to analysis application 260 of Fig. 4. However,
2 applications 314 and 316 different from applications 260 and 258 in that
3 applications 314 and 316 are part of the same game device 300. So, any
4 commands or data sent between application 314 and application 316 are sent
5 internal to game device 300 rather than between two devices. Additionally,
6 analysis control and feedback application 314 presents its user interface by way of
7 display device 310.

8 The analysis control and feedback application (e.g., application 258 of Fig.
9 4 or application 314 of Fig. 5) can present feedback to the user using any of a
10 variety of user interfaces, including graphical user interfaces. Fig. 6 illustrates an
11 example user interface that may be presented by the analysis control and feedback
12 application. In the example of Fig. 6, the user interface display 340 includes a
13 timeline window 342, an events window 344, and a frame window 346.

14 Timeline window 342 displays the captured timing information graphically
15 using a timeline. This captured timing information can include, for example,
16 when particular CPU or GPU events occurred.

17 Events window 344 contains a listing or grid of all of the events that have
18 been captured along with all of the data regarding the frame that has been
19 calculated as part of the analysis of the frame (e.g., by analysis application 260 of
20 Fig. 4 or analysis application 316 of Fig. 5).

21 Frame window 346 displays information regarding the display of the frame,
22 and may include images showing how the frame appears at different points during
23 its being drawn. Windows 342, 344, and 346 are discussed in additional detail
24 below.

1 It should be noted that although all three windows 342, 344, and 346 are
2 illustrated in Fig. 6, alternatively only one or more of the windows may be
3 displayed. For example, user interface display 340 may display only window 342,
4 only window 344, only window 346, windows 342 and 344 but not window 346,
5 windows 344 and 346 but not window 342, or windows 342 and 346 but not
6 window 344.

7 Windows 342, 344, and 346 are illustrated in Fig. 6 with particular
8 positions relative to one another. Alternatively, one or more of windows 342, 344,
9 and 346 may be positioned elsewhere. Additionally, windows 342, 344, and 346
10 are illustrated in Fig. 6 as being adjacent to one another. Alternatively, one or
11 more of windows 342, 344, and 346 may be positioned away from (not adjacent
12 to) one or more of the other windows 342, 344, and 346.

13 Fig. 7 illustrates another example user interface 360 that may be presented
14 by the analysis control and feedback application. User interface 360 illustrates a
15 more detailed version of user interface 340 of Fig. 6. User interface 360 includes
16 a timeline window 362, an events window 364, and a frame window 366.
17 Timeline window 362 includes a timeline 368 and an associated time scale 370.
18 Time scale 370 refers to the drawing of a single frame (e.g., 0 ms refers to the
19 beginning of the drawing of the frame). A time cursor may also be displayed in
20 timeline window 362, the time cursor identifying a currently selected event on
21 timeline 368. Timeline 368 includes multiple horizontal bars, each bar
22 corresponding to a particular event in the drawing of the frame. Given the
23 location of a particular bar and time scale 370, the time during drawing of the
24 frame at which the event occurred can be determined by the user. Additionally,
25

1 the width of each horizontal bar signifies how long it took for the event to be
2 drawn.

3 An event as discussed herein can be different things in different
4 embodiments. For example, in certain embodiments there may be particular
5 commands that constitute events, such as those listed in Table I above. In other
6 embodiments, every command submitted to the graphics processing unit may be
7 an event. In other embodiments, the events may be identified by the game
8 designer. For example, as discussed above the game designer may design a game
9 to specify certain information when the game is submitting commands to the
10 graphics processing unit (e.g., identifying what part of the frame is being drawn,
11 such as the character's left arm, the character's right arm, the character's weapon,
12 the background terrain, etc.). Such information can be used as events (e.g., each
13 part of the frame being drawn can be a separate event).

14 Timeline window 362 displays the events that are listed in events window
15 364. These events can be hierarchical in nature. For example, a MainCharacter
16 event may contain other events (referred to as children events) for DrawLeftArm,
17 DrawHead, DrawWeapon, and so forth. When one event contains other events, it
18 is displayed with a small plus sign next to it in events window 364. When the plus
19 sign is clicked, the event is expanded so that its children events are displayed.
20 When this happens, the event display in timeline window 362 changes. The
21 horizontal bar that represented the parent event is now replaced with one or more
22 horizontal bars that represent the child events. These horizontal bars may be the
23 same color as the parent event or a different color. the space covered by the
24 children events on the timeline does not exceed the space covered by the parent
25 event.

1 All of the bars in timeline 368 may be the same color, or different colors
2 may be used to signify different events. For example, clear commands may be one
3 color, draw commands may be another color, swap commands may be yet another
4 color, and so forth.

5 Additionally, more specific information may be available to the user by
6 moving a cursor (e.g., an arrow or other on-screen cursor) over a part of timeline
7 368. The cursor can be moved using any of a variety of devices, such as a mouse,
8 game controller, keyboard, track pad, trackball, and so forth. When the cursor is
9 stationary over an event for greater than a threshold amount of time (e.g., one or
10 two seconds), more specific information can be displayed to the user.
11 Alternatively, when the event is "selected" (such as by depressing a mouse button
12 or game controller button when the cursor is over the event) the more specific
13 information can be displayed to the user. This more specific information can take
14 a variety of different forms, such as the specific start and end times for the event
15 (e.g., in nanoseconds).

16 Timeline window 362 can include a timeline for the CPU, a timeline for the
17 GPU, or timelines for both the CPU and the GPU. Fig. 8 illustrates an example
18 timeline window 380 which could be displayed as timeline window 362 of Fig. 7.
19 Timeline window 380 includes a CPU timeline 382 on which CPU events are
20 displayed and a GPU timeline 384 on which GPU events are displayed. In many
21 instances, a CPU event has a corresponding GPU event, and in such situations
22 when either the CPU event or the GPU event is selected (e.g., by depressing a
23 button on a cursor control device when the cursor is over the event, or moving a
24 directional key (e.g., the + and - keys) of a keyboard), an arrow is displayed to
25 link the CPU and GPU events.

1 When a CPU event is selected, the head of the arrow that points to the CPU
2 event is solid and the head of the arrow that points to the corresponding GPU
3 event is hollow. When a CPU event is selected, the arrow keys cause the CPU
4 cursor to move linearly along CPU timeline 382. The down-arrow key moves the
5 solid headed arrow down to the GPU event and subsequent arrow-key presses
6 move the cursor linearly along GPU timeline 384. Pressing the up-arrow key
7 moves the solid arrow head back up to CPU timeline 382 and the corresponding
8 CPU event is selected. When a CPU event is selected, the other arrow points to
9 the corresponding GPU event. Likewise, when a GPU event is selected, the other
10 arrow points to the corresponding CPU event.

11 One or more graphs 386 may also be included in timeline window 380.
12 Different information may be graphed, and in certain embodiments the user is able
13 to select one of the columns from the events window (e.g., window 364 of Fig. 7,
14 as discussed in more detail below). The data from the selected column is then
15 plotted against time in graph 386. In the example graph 386, time is along the
16 horizontal axis and the data from the selected column is along the vertical axis.
17 Different columns can be selected from pull-down menu 388.

18 Returning to Fig. 7, events window 364 includes a table with a columnar
19 listing of all of the events for the frame that have been captured along with all of
20 the data regarding the frame that has been calculated as part of the analysis of the
21 frame (e.g., by analysis application 260 of Fig. 4 or analysis application 316 of
22 Fig. 5). Each row in the table represents an event. A CPU-only event will contain
23 values in the CPU Start and CPU Duration columns. An event that also occurs on
24 the GPU will have valid values in the GPU Start and Duration columns. If there is
25 no timing information available, none of these four columns will have data. The

1 remaining columns are calculated by the analysis application and will be empty if
2 no analysis has been performed. Fig. 9 illustrates another example events window
3 400 with additional data.

4 The ordering of the events listed in events window 364 is ordinarily
5 determined by the CPU Start Time column. Clicking on another column will
6 cause the events to be listed in ascending order for that column. Clicking again on
7 the header will cause the events to be listed in descending order for that column.
8 An option is provided in the context menu (e.g., activated by right-clicking in the
9 window) to restore the event sorting to its default value. Each column can be
10 resized by dragging the edge of the column header with a cursor control device.
11 The headers remain in place vertically as the events are scrolled up or down but
12 they track the horizontal movement of the table as it is scrolled from side-to-side.
13 Right clicking on the table or the columns brings up a context menu with
14 additional options for resizing or other operations.

15 One example of such an operation is to export the entire table to a file that
16 can be loaded into a spreadsheet or other program for further analysis. Another
17 example of such an operation is to specify one or more columns that will be
18 graphed in the timeline window along with the GPU event data (e.g., as graph 386
19 of Fig. 8). Another example of such an operation is display of a column chooser
20 to assist the user in managing the display of the columns. User-selection of this
21 operation causes a temporary window (not shown) to be displayed that will allow
22 a set of columns and their order to be chosen. Columns that are not chosen will
23 not be displayed in window 364. One or more predetermined column
24 configurations for particular tasks (such as analyzing fill rate) may be presented to
25 the user as part of the column chooser.

1 One or more rows of the table can be highlighted. The highlighted row(s)
2 represent the event(s) which correspond to the event(s) that are represented by the
3 currently selected horizontal bar of timeline 368. When a new horizontal bar(s) of
4 timeline 368 is selected, then a new row(s) of the table is highlighted (the new
5 row(s) representing the same event(s) as the newly selected horizontal bar(s)).
6 Similarly, if a new row(s) of the table is selected and highlighted, then a new
7 horizontal bar(s) of timeline 368 is selected (the new horizontal bar(s) representing
8 the same event(s) as the newly selected row(s).

9 The last row in the table is a summary row. The summary row contains
10 calculated values based on all of the other events in the table. Each column has a
11 summary type associated with it that is one of a maximum value, a minimum
12 value, an average value, or a sum value. The summary value is calculated based
13 on the summary type and the data in the column.

14 The events window will also support a hierarchical display of events. An
15 application can specify user-defined events in a hierarchical way. When these
16 events encompass other events, the user will have an option to display the user-
17 event, the CPU and GPU events, and/or other hierarchical events that comprise it.
18 For example, a game may specify a Start time for their physics engine and an end
19 time. If no graphics events occurred during this time, the physics engine event
20 will show up as a discrete CPU-only event in the events window. A game may
21 also specify a start time for rendering the main character in the game and also an
22 end time. In between these times are all the rendering calls for the main character.
23 This event can be displayed in the events window with an expansion icon next to
24 it. If the user clicks on this expansion icon, the events window will display all of
25 the graphics events that occurred during the user event. When the user event is

1 displayed, its columns will contain a summary of the information contained by the
2 events under it in the hierarchy based on the summary type specified for each
3 column.

4 By way of example, events window 400 of Fig. 9 shows a hierarchical
5 event Bear Mesh 0 along with its children events (DrawIndexedVertices,
6 KickPushBuffer, KickPushBuffer, DrawIndexedVertices, and
7 DrawIndexedVertices). Other hierarchical events are also shown, such as Bear
8 Mesh 1, although its expansion icon has not been selected so its children events
9 are not displayed.

10 A variety of different columns can be included in events window 364.
11 Table II lists examples of such columns. Reference is made in Table II, as well as
12 elsewhere in this description, to the "push buffer" or "pushbuffer". The push
13 buffer or pushbuffer refers to the commands and their associated data that are
14 submitted to the graphics processing unit to be drawn. These commands and their
15 associated data are captured, as discussed above.

Table II

Column	Description
Back-end Time	Total back-end time that the rendering primitive took, as measured at the Z-cull stage of the graphics hardware pipeline. A hardware back-end flush is done between primitives to measure each primitive independently from one another. Because the flush eliminates overlaps in the pipeline between primitives, the sum of the back-end times may be more than the overall back-end render time.
Pre-Occlusion Cull Pixel Count	The number of pixels that would have been rendered if depth, stencil, and alpha comparison functions were ALWAYS. When multi-sampling, this is the post-multisample-expansion count.
Post-Occlusion Cull Pixel Count	The number of pixels that are actually rendered, taking into account depth, stencil, and alpha occlusion. When multi-sampling, this is the post-multisample-expansion count.
Pixels Occlusion Culled	Percent of pixels that are occlusion culled.
Effective Fill Rate	Effective pixel fill rate, relative to the back-end time.
Ideal Fill Rate	Ideal pixel fill rate, relative to the back-end time. This is measured using all of the current state with the current render target, but using large quads that are rendered such that Z and stencil tests always pass, and with all textures forced to be opaque and 1x1 (or 1x1x1) texels in size. This is the fill-rate that could be achieved with the current primitive in the current state if the current Z compression was 100%, there was 100% utilization of the 4 pixel pipelines, and the texture reads had no performance impact.
Vertex Count	Vertex count.
Triangle Count	Triangle Count.
Effective Triangle Rate	Effective triangle fill rate, relative to the back-end time.
Vertex Shader Cost	Measured speed of the current vertex shader program (or fixed-function T&L), in GPU cycles.
Back-end Time With 1-Pixel Textures	The same measurement as the back-end time, but done with all textures set to opaque, 1x1 (or 1x1x1) texel size.
% Texture Bound	Measures how much of the rendered primitive time is attributable to texture fetches, as computed using the back-end times for 1x1 and normal sized textures.
Back-end Time	The same measurement as the back-end time, but done with

1 With 0-Pixel Viewport	the viewport forced to a zero-pixel size.
2 % Fill Bound	Measures how much of the rendered primitive time is attributable to fills, as computed using the back-end times for zero-pixel and normal sized viewports.
3 Z-Compressed Packets	The percentage of Z packets that are compressed after the primitive is rendered. Note that this is not the effective compression of the Z-buffer (after a Z clear, the percentage of Z-compressed packets will be 100%, but the effective compression is 87.5% due to 8-to-1 compression).
4 Push-buffer Inline Data	The number of bytes that are written into the push buffer to handle the command for the primitive rendering. For DrawIndexedVertices, this is effectively the amount of index data that has to be copied to the push buffer.
5 Push-buffer Setup Data	The number of bytes of state change commands in the push buffer that preceded the rendered primitive.
6 Push Buffer Event	Where the event occurs chronologically in drawing of the frame.

12 Returning to Fig. 7, frame window 366 displays information regarding the
 13 display of the frame. This information regarding the frame may be a display of the
 14 frame as drawn, or different aspects of the frame (e.g., a depth buffer or wireframe
 15 view, shader programs used, warnings, and so forth as discussed below).
 16 Depending on the type of information being displayed, the information may be for
 17 the entire frame or the information may be displayed for different chronological
 18 points relative to the drawing of the frame. For example, if the information
 19 comprises warnings, then the information is for the entire frame. By way of
 20 another example, the information displayed in frame window 366 may show
 21 information regarding display of the frame as it appeared after being fully drawn,
 22 as it appeared after being only 25% drawn, as it appeared immediately after a
 23 particular event was finished being drawn, and so forth. Which chronological
 24 point is displayed in frame window 366 can vary based on which event in events
 25 window 364 is selected. Whenever an event in events window 364 is selected,

1 frame window 366 displays information regarding display of the frame as it
2 appeared immediately after the selected event was finished being drawn.

3 Different windows can be displayed as the frame window 366. Examples
4 of such windows that can be displayed as the frame window 366 include an
5 images window (examples of which are illustrated in Figs. 10-15), a call stack
6 window (an example of which is illustrated in Fig. 16), a warnings window (an
7 example of which is illustrated in Fig. 17), a pushbuffer window (an example of
8 which is illustrated in Fig. 18), a summary window (an example of which is
9 illustrated in Fig. 19), a shaders window (an example of which is illustrated in Fig.
10 20), a renderstates window (an example of which is illustrated in Fig. 21), a
11 texturestates window (an example of which is illustrated in Fig. 22), and an other
12 state window (an example of which is illustrated in Fig. 23). Frame window 366
13 displays a set of tabs 410 which can be selected by the user to determine the
14 window to be displayed as frame window 366.

15 Fig. 10 illustrates an example of an images window 420 that can be
16 displayed as frame window 366 of Fig. 7. The images window is used to display
17 the many surfaces that are used by the game. Images window 420 of Fig. 7
18 illustrates a render target view 422. Different views for images window 420 can
19 be selected from pull down menu 424. Examples of different views that can be
20 selected for images window 420 include a render target view (an example of
21 which is illustrated in Fig. 10), a depth buffer view (an example of which is
22 illustrated in Fig. 11), a wireframe view (an example of which is illustrated in Fig.
23 12), a texture view (an example of which is illustrated in Fig. 13), an overdraw
24 view (an example of which is illustrated in Fig. 14), and a fill rate view (an
25 example of which is illustrated in Fig. 15).

1 The surfaces can be displayed within the images window 420 at various
2 sizes (e.g., 100%, 200%, 400%, 50%, 25%, fit to window, etc.).

3 When the mouse is moved over a particular place on a surface displayed in
4 the images window, the actual values at that pixel will be displayed in a status bar.
5 An example status bar 426 is illustrated in Fig. 10.

6 For one or more of the views of the images window, such as the render
7 target view, a particular pixel or group of pixels can be selected by the user. This
8 selection can be made, for example, by the user moving a cursor over a pixel and
9 activating the pointing device being used to control the cursor (e.g., clicking on the
10 mouse button). By way of another example, a user could click and drag a
11 rectangle to select a group of pixels. When such a pixel(s) is selected, all of the
12 events in the events window (view 364 of Fig. 7) that render to that pixel will be
13 highlighted.

14 The render target view 422 shows the color and alpha values held by the
15 frame buffer. These values can correspond to the final image that is the frame, or
16 alternatively may be some point during drawing of the frame. For example, as
17 discussed above, only the commands for drawing a particular part of the scene
18 may be submitted to the graphics processing unit, and the resultant frame can be
19 displayed using the render target view.

20 In certain embodiments, the render target view 422 in conjunction with the
21 timeline in timeline window 362 or the events in events window 364 can allow a
22 user to "step through" the drawing of a frame event by event. For example, the
23 user can select the first event (the first row) in the grid of window 364 and have
24 the resultant view after only that first event is drawn by the graphics processing
25 unit displayed as the render target view 422. The user can then select the second

1 event (the second row) in the grid of window 364 and have the resultant view after
2 both the first and second events are drawn by the graphics processing unit
3 displayed as the render target view 422. This process can continue, allowing the
4 user to slowly see how the frame is drawn and to easily identify any drawing
5 problems associated with each event.

6 Fig. 11 illustrates an example of the depth buffer view of images window
7 420. The depth buffer view shows the depth buffer associated with the render
8 target. This buffer holds depth values for each rendered pixel. The depth values
9 are displayed in the window with the closest pixels colored white and the farthest
10 pixels colored black with the rest of the pixels some shade in between based on
11 their depth. Alternatively, other color schemes can be used.

12 Fig. 12 illustrates an example of the wireframe view of images window
13 420. The wireframe view shows the outline of each triangle rendered in the frame.
14 The outline is shown in the same color as the outside border of pixels in the
15 triangle. This view allows a user to see how many triangles there are in a
16 particular scene and how dense they are. It can also show problems with
17 geometry.

18 Fig. 13 illustrates an example of the texture view of images window 420.
19 In certain embodiments, the game device may support multiple different textures
20 for a given pixel. A different texture view can be displayed for each of these
21 multiple different textures, with the view for a particular texture showing the
22 contents of that texture. A combination texture view can also be displayed that
23 shows two or more of these multiple different textures concurrently. Fig. 13
24 illustrates an example of an all textures window view that displays the contents of
25 four textures at once (texture 428, texture 430, texture 432, and texture 434). It

1 should be noted that textures can be different sizes, and in the example of Fig. 13
2 all four textures are displayed using the same scale (thus, texture 434 appears
3 much smaller than textures 428, 430, and 432).

4 Fig. 14 illustrates an example of the overdraw view of images window 420.
5 The overdraw view shows graphically how many times each pixel on the screen is
6 drawn. A pixel may be drawn multiple times because of transparency or because
7 one object occludes another object that has already been drawn. In the overdraw
8 view, pixels are shown more brightly each time they are touched during the scene.
9 Black pixels have not been drawn at all. The whiter the pixel, the more times it
10 has been touched. Alternatively, other color schemes can be used. This view is
11 useful to users so that they can see which pixels are being drawn several times
12 during the scene. Excessive overdraw is something that game developers may try
13 to reduce in order to make their games run faster.

14 Fig. 15 illustrates an example of the fill rate view of images window 420.
15 The fill rate view shows how fast the GPU was running when each pixel was
16 drawn. For example, if the game is using a complicated vertex or pixel shader
17 program, each pixel takes longer to draw and the fillrate that the GPU is able to
18 sustain is less. In this view, the more expensive pixels (those taking longer to
19 draw) are brighter. This lets the user know which pixels are being drawn faster
20 than other pixels.

21 Fig. 16 illustrates an example of a call stack window 450 that can be
22 displayed as frame window 366 of Fig. 7. Stack trace data can be captured, as
23 discussed above. The stack trace data captured is displayed in call stack window
24 450. Call stack window 450 scrolls to show the current event selected in either the
25 timeline window (e.g., window 362 of Fig. 7) or events window (e.g., window 364

1 of Fig. 7). The user can specify the directory where a symbol file is located (e.g.,
2 using directory entry box 452). The information in the symbol file can then be
3 used to identify the routine names, file names, and line numbers for the captured
4 stack trace, which in turn can be displayed to the user.

5 In the example call stack window 450, the stack trace data displayed to the
6 user shows the sequence of calls that resulted in the listed event being sent as a
7 command to the GPU. Each line in the display corresponds to a function call in
8 the program. The first function listed shows the function that submitted a
9 command to the GPU that resulted in the listed event. The next line shows the
10 function that called the first function, etc. The last function listed for each event is
11 the top-level function in the game. The filename and line number where each
12 function is located in the game source code is also displayed.

13 Fig. 17 illustrates an example of a warning window 460 that can be
14 displayed as frame window 366 of Fig. 7. As discussed above, as data for a frame
15 is analyzed, violations of one or more recommendations or rules may be detected.
16 Warnings regarding such violations can be displayed in warning window 460. In
17 the example warning window 460, a list of events in the frame and warnings (if
18 any) for each event are shown in a table format. In certain embodiments, an icon
19 (such as a "!") may be placed next to events in the events (e.g., window 364 of Fig.
20 7) to notify the user that there is a warning associated with that event. The user
21 can then select this icon (e.g., by clicking on the icon using a cursor and cursor
22 control device), in response to which warning window 460 is displayed as frame
23 window 366 and the event and warning(s) in window 460 associated with that icon
24 scrolled to (if necessary) and highlighted in window 460. Warning window 460
25

1 scrolls to show the current event selected in either the timeline window (e.g.,
2 window 362 of Fig. 7) or events window (e.g., window 364 of Fig. 7).

3 Each warning has a priority associated with it from 1 to 3. Priority 1
4 warnings are serious problems that may have a significant impact on the time it
5 takes the game to draw a frame. Priority 2 warnings are less important but still
6 important to consider when performance-tuning a game. Priority 3 warnings are
7 suggestions that are good to follow but may not have a large impact on the overall
8 performance. The user can use check boxes 462 to choose which warnings should
9 be displayed according to their priority.

10 Fig. 18 illustrates an example of a pushbuffer window 464 that can be
11 displayed as frame window 366 of Fig. 7. The pushbuffer window shows a
12 pseudo-disassembly of the captured commands sent to the graphics processing unit
13 for an event. The pushbuffer window shows all of the commands that were
14 captured in for the current frame. In the window, each event that unnecessarily
15 sets state to the same value that it previously had is marked as “Redundant”. The
16 pushbuffer window scrolls to show the current event selected in either the timeline
17 window (e.g., window 362 of Fig. 7) or events window (e.g., window 364 of Fig.
18 7). Clicking on an event in the pushbuffer window will set a time cursor to the
19 corresponding event in the timeline window (e.g., window 362 of Fig. 7) and the
20 other windows (e.g., the events window or other windows that can be displayed as
21 frame window 366) will update to the newly selected event.

22 Fig. 19 illustrates an example of a summary window 468 that can be
23 displayed as frame window 366 of Fig. 7. The summary window provides
24 information that applies to the entire frame. Examples of this type of information
25 include the frame rate, a summary of state changes, overall memory usage, etc.

1 Fig. 20 illustrates an example of a shaders window 472 that can be
2 displayed as frame window 366 of Fig. 7. The shaders window shows the code of
3 the current pixel or vertex shader programs. A pull-down menu 474 can be used
4 to select between the current pixel shader program and the current vertex shader
5 program. The pixel and/or vertex shader code could be pseudo-code, source code,
6 etc. A copy option 476 can also be selected by the user to have the displayed code
7 of the current pixel and/or vertex shader programs copied to the clipboard for
8 subsequent use with another program (e.g., word processing program or text
9 editor).

10 Fig. 21 illustrates an example of a renderstates window 480 that can be
11 displayed as frame window 366 of Fig. 7. The renderstates window 480 shows the
12 current GPU state for render states. The states and their values for the currently
13 selected event is displayed in window 480. When a new event is selected by the
14 user by clicking in the Timeline window or the Events window, the renderstate
15 values displayed change to reflect the state of the GPU at the time that event was
16 executed.

17 Fig. 22 illustrates an example of a texturestates window 484 that can be
18 displayed as frame window 366 of Fig. 7. The texturestates window 484 shows
19 the current GPU state for texture states. The states and their values for the
20 currently selected event is displayed in window 484. When a new event is
21 selected by the user by clicking in the Timeline window or the Events window, the
22 texturestate values displayed change to reflect the state of the GPU at the time that
23 event was executed.

24 Fig. 23 illustrates an example of an otherstates window 488 that can be
25 displayed as frame window 366 of Fig. 7. The otherstates window 488 shows the

1 current GPU state for other states (other than render states and texture states). The
2 states and their values for the currently selected event are displayed in window
3 480. When a new event is selected by the user by clicking in the Timeline
4 window or the Events window, the state values displayed change to reflect the
5 state of the GPU at the time that event was executed.

6 Fig. 24 illustrates an example pixel history window 500 that can also be
7 displayed by the analysis control and feedback application. The pixel history
8 window 500 is displayed as a separate window (e.g., another window in addition
9 to windows 342, 344, and 346 of Fig. 6). Pixel history window 500 lists each
10 event that affects a selected pixel during the frame up to the currently selected
11 event. The pixel can be selected by the user in different manners. In one
12 implementation, when the user clicks on the right mouse button (or other button of
13 a cursor control device), the pixel that the cursor is over is selected and pixel
14 history window 500 is opened.

15 Pixel history window 500 lists each event that affects the selected pixel
16 during the frame, or alternatively only up to a particular point in the frame (e.g., as
17 identified by the user). Additional information regarding each event as it affects
18 the selected pixel is also displayed. The specific information displayed can vary
19 based on the event. In the illustrated example of pixel history window 500, initial
20 frame buffer values 502 are displayed. The frame then starts with a clear
21 command, which is shown as clear event 504. Each subsequent event 506
22 affecting the selected pixel is also identified in window 500, as well as additional
23 information regarding those events. For example, the pixel shader output color
24 and frame buffer color after blending may can be identified.

Pixel history window 500 may also include links to the event, a pixel shader debugger, a vertex shader debugger, and a mesh debugger. If the link to the event is selected by the user, then that event is made the current event in the user interface (e.g., windows 362, 364, and 366 of Fig. 7) and the other windows are updated accordingly. If the pixel shader debugger link is selected by the user, then a pixel shader debugger window is opened to display information regarding the pixel shader used for the selected pixel. If the mesh debugger link is selected by the user, then a mesh debugger window is opened to display information regarding the triangle currently being drawn (as of the current event) that the selected pixel is part of.

Fig. 25 illustrates an example pixel shader debugger window 520 that can also be displayed by the analysis control and feedback application. Pixel shader debugger window 520 is displayed in place of pixel history window 500 of Fig. 24. Clicking the “Back” button on the pixel shader debugger window causes the pixel history window to be displayed again. Pixel shader debugger window 520 shows the pixel shader program that executed for the selected pixel. The input and output register values are displayed for each instruction in the pixel shader debugger program. Pixel shader debugger window 520 allows, for example, a user to see exactly what happened when this program ran for a particular pixel and why the result might have been different than what was expected.

Pixel shader debugger window 520 also includes a vertex portion 522 that includes information regarding the vertices of the triangle currently being drawn (as part of the current event) that the selected pixel is part of. The vertex shader debugger window can be displayed in response to a user selection of such an

option in window 520, such as user selection of one of the vertices identified in vertex portion 522.

Fig. 26 illustrates an example vertex shader debugger window 540 that can also be displayed by the analysis control and feedback application. Vertex shader debugger window 540 is displayed in place of pixel history window 500 of Fig. 24. The vertex shader debugger window 540 is similar to the pixel shader debugger window 520 of Fig. 25. Vertex shader debugger window 540 displays the input and output register values for each instruction in the vertex shader program for the triangle currently being rendered.

Fig. 27 illustrates an example mesh debugger window 560 that can also be displayed by the analysis control and feedback application. Mesh debugger window 560 is displayed in place of pixel history window 500 of Fig. 24. Clicking the “Back” button on the mesh debugger window causes the pixel history window to be displayed again. The mesh debugger window 560 can be launched from the pixel history window 500 of Fig. 24 as discussed above, or can be launched directly from the events window 364 of Fig. 7 (e.g., by the user clicking on the right mouse button (or other button of a cursor control device) when the cursor is over a particular event). The mesh debugger window 560 displays information about a single mesh of the frame. A mesh contains geometry information for triangles that make up an object. The mesh debugger window 560 displays two wireframe views of the mesh (the first wireframe view 562 shows the edges of every triangle in the mesh, while the second wireframe view 564 shows the edges of every triangle in the mesh relative to the screen position) and a table that shows the attribute values for each vertex in the mesh as well as the output of the vertex shader program for each vertex in the mesh. Additionally, the vertex

1 shader debugger window 540 of Fig. 26 can be launched by clicking on the index
2 value for a particular vertex.

3 It should be noted that, with reference to Figs. 6 – 27 above, as different
4 views and/or windows are selected, the selected pixel(s) and/or events remain the
5 same. For example, a user could select a pixel in the render target view of images
6 window 420 (of Fig. 10), and then select the shaders window 472 (of Fig. 20) to
7 have the pixel and/or vertex shader program for that selected pixel displayed.

8 Although the discussions above primarily discuss facilitating performance
9 analysis for graphics processing, the techniques described herein can similarly be
10 used to facilitate performance analysis for other types of processors, including
11 general purpose processors and specific purpose (e.g., dedicated) processors. The
12 techniques described herein can be applied to virtually any processor that accepts a
13 command stream and can be used to provide information regarding processing of
14 the command stream. For example, the techniques described herein can be used to
15 facilitate performance analysis for general purpose processors (e.g., CPUs).
16 Correspondingly, the techniques described herein can be used to facilitate
17 performance analysis for processors on other computing devices other than game
18 devices.

19 Depending on the type of application being executed on the computing
20 device and the type of processor, the capture of commands sent to the processor
21 can vary. For example, if the application is a graphics or video application, then
22 the application may still send commands to the processor indicating the beginning
23 and ending of frames. Thus, the commands to be captured can be identified on a
24 frame-by-frame basis. In other examples, the application may not issue any such
25 indications of the beginning and ending of frames. In such situations, the

1 beginning and ending points for capture would be identified in different manners.
2 For example, a separate command(s) may be embedded in the application that
3 indicate the beginning and ending of points for capture, or the capture may simply
4 begin as soon as the request to capture is received, and then end after some amount
5 of time (e.g., the request may indicate how long the capture should continue for,
6 such as a length of time in milliseconds or seconds, or a number of commands to
7 be captured, etc.).

8 Capturing commands for a general purpose processor can permit a great
9 deal of information regarding processing of those commands by the processor to
10 be returned to the application designer. For example, using the techniques
11 described herein precise measures of how long processing of particular commands
12 or groups of commands took can be made. By way of another example,
13 instruction traces identifying the specific instructions executed in the captured
14 portion can be readily determined.

15 Fig. 28 illustrates an example of a general computer environment 600,
16 which can be used to implement the techniques described herein. The computer
17 environment 600 is only one example of a computing environment and is not
18 intended to suggest any limitation as to the scope of use or functionality of the
19 computer and network architectures. Neither should the computer environment
20 600 be interpreted as having any dependency or requirement relating to any one or
21 combination of components illustrated in the example computer environment 600.

22 Computer environment 600 includes a general-purpose computing device in
23 the form of a computer 602. Computer 602 can be, for example, game device 152
24 or computing device 154 of Fig. 2, game device 200 of Fig. 3, game device 242 or
25 computing device 244 of Fig. 4, or game device 300 of Fig. 5, and may implement

1 process 100 of Fig. 1. The components of computer 602 can include, but are not
2 limited to, one or more processors or processing units 604, a system memory 606,
3 and a system bus 608 that couples various system components including the
4 processor 604 to the system memory 606.

5 The system bus 608 represents one or more of any of several types of bus
6 structures, including a memory bus or memory controller, a peripheral bus, an
7 accelerated graphics port, and a processor or local bus using any of a variety of
8 bus architectures. By way of example, such architectures can include an Industry
9 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an
10 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)
11 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a
12 Mezzanine bus.

13 Computer 602 typically includes a variety of computer readable media.
14 Such media can be any available media that is accessible by computer 602 and
15 includes both volatile and non-volatile media, removable and non-removable
16 media.

17 The system memory 606 includes computer readable media in the form of
18 volatile memory, such as random access memory (RAM) 610, and/or non-volatile
19 memory, such as read only memory (ROM) 612. A basic input/output system
20 (BIOS) 614, containing the basic routines that help to transfer information
21 between elements within computer 602, such as during start-up, is stored in ROM
22 612. RAM 610 typically contains data and/or program modules that are
23 immediately accessible to and/or presently operated on by the processing unit 604.

24 Computer 602 may also include other removable/non-removable,
25 volatile/non-volatile computer storage media. By way of example, Fig. 28

1 illustrates a hard disk drive 616 for reading from and writing to a non-removable,
2 non-volatile magnetic media (not shown), a magnetic disk drive 618 for reading
3 from and writing to a removable, non-volatile magnetic disk 620 (e.g., a “floppy
4 disk”), and an optical disk drive 622 for reading from and/or writing to a
5 removable, non-volatile optical disk 624 such as a CD-ROM, DVD-ROM, or other
6 optical media. The hard disk drive 616, magnetic disk drive 618, and optical disk
7 drive 622 are each connected to the system bus 608 by one or more data media
8 interfaces 626. Alternatively, the hard disk drive 616, magnetic disk drive 618,
9 and optical disk drive 622 can be connected to the system bus 608 by one or more
10 interfaces (not shown).

11 The disk drives and their associated computer-readable media provide non-
12 volatile storage of computer readable instructions, data structures, program
13 modules, and other data for computer 602. Although the example illustrates a hard
14 disk 616, a removable magnetic disk 620, and a removable optical disk 624, it is to
15 be appreciated that other types of computer readable media which can store data
16 that is accessible by a computer, such as magnetic cassettes or other magnetic
17 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
18 other optical storage, random access memories (RAM), read only memories
19 (ROM), electrically erasable programmable read-only memory (EEPROM), and
20 the like, can also be utilized to implement the example computing system and
21 environment.

22 Any number of program modules can be stored on the hard disk 616,
23 magnetic disk 620, optical disk 624, ROM 612, and/or RAM 610, including by
24 way of example, an operating system 626, one or more application programs 628,
25 other program modules 630, and program data 632. Each of such operating

1 system 626, one or more application programs 628, other program modules 630,
2 and program data 632 (or some combination thereof) may implement all or part of
3 the resident components that support the distributed file system.

4 A user can enter commands and information into computer 602 via input
5 devices such as a keyboard 634 and a pointing device 636 (e.g., a "mouse").
6 Other input devices 638 (not shown specifically) may include a microphone,
7 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
8 other input devices are connected to the processing unit 604 via input/output
9 interfaces 640 that are coupled to the system bus 608, but may be connected by
10 other interface and bus structures, such as a parallel port, game port, or a universal
11 serial bus (USB).

12 A monitor 642 or other type of display device can also be connected to the
13 system bus 608 via an interface, such as a video adapter 644. In addition to the
14 monitor 642, other output peripheral devices can include components such as
15 speakers (not shown) and a printer 646 which can be connected to computer 602
16 via the input/output interfaces 640.

17 Computer 602 can operate in a networked environment using logical
18 connections to one or more remote computers, such as a remote computing device
19 648. By way of example, the remote computing device 648 can be a personal
20 computer, portable computer, a server, a router, a network computer, a peer device
21 or other common network node, and the like. The remote computing device 648 is
22 illustrated as a portable computer that can include many or all of the elements and
23 features described herein relative to computer 602.

24 Logical connections between computer 602 and the remote computer 648
25 are depicted as a local area network (LAN) 650 and a general wide area network

1 (WAN) 652. Such networking environments are commonplace in offices,
2 enterprise-wide computer networks, intranets, and the Internet.

3 When implemented in a LAN networking environment, the computer 602 is
4 connected to a local network 650 via a network interface or adapter 654. When
5 implemented in a WAN networking environment, the computer 602 typically
6 includes a modem 656 or other means for establishing communications over the
7 wide network 652. The modem 656, which can be internal or external to computer
8 602, can be connected to the system bus 608 via the input/output interfaces 640 or
9 other appropriate mechanisms. It is to be appreciated that the illustrated network
10 connections are examples and that other means of establishing communication
11 link(s) between the computers 602 and 648 can be employed.

12 In a networked environment, such as that illustrated with computing
13 environment 600, program modules depicted relative to the computer 602, or
14 portions thereof, may be stored in a remote memory storage device. By way of
15 example, remote application programs 658 reside on a memory device of remote
16 computer 648. For purposes of illustration, application programs and other
17 executable program components such as the operating system are illustrated herein
18 as discrete blocks, although it is recognized that such programs and components
19 reside at various times in different storage components of the computing device
20 602, and are executed by the data processor(s) of the computer.

21 Various modules and techniques may be described herein in the general
22 context of computer-executable instructions, such as program modules, executed
23 by one or more computers or other devices. Generally, program modules include
24 routines, programs, objects, components, data structures, etc. that perform
25 particular tasks or implement particular abstract data types. Typically, the

1 functionality of the program modules may be combined or distributed as desired in
2 various embodiments.

3 An implementation of these modules and techniques may be stored on or
4 transmitted across some form of computer readable media. Computer readable
5 media can be any available media that can be accessed by a computer. By way of
6 example, and not limitation, computer readable media may comprise "computer
7 storage media" and "communications media."

8 "Computer storage media" includes volatile and non-volatile, removable
9 and non-removable media implemented in any method or technology for storage
10 of information such as computer readable instructions, data structures, program
11 modules, or other data. Computer storage media includes, but is not limited to,
12 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
13 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
14 tape, magnetic disk storage or other magnetic storage devices, or any other
15 medium which can be used to store the desired information and which can be
16 accessed by a computer.

17 "Communication media" typically embodies computer readable
18 instructions, data structures, program modules, or other data in a modulated data
19 signal, such as carrier wave or other transport mechanism. Communication media
20 also includes any information delivery media. The term "modulated data signal"
21 means a signal that has one or more of its characteristics set or changed in such a
22 manner as to encode information in the signal. By way of example, and not
23 limitation, communication media includes wired media such as a wired network or
24 direct-wired connection, and wireless media such as acoustic, RF, infrared, and
25

1 other wireless media. Combinations of any of the above are also included within
2 the scope of computer readable media.

3 One or more flowcharts are described herein and illustrated in the
4 accompanying Figures. The ordering of acts in these flowchart(s) are examples
5 only – these orderings can be changed so that the acts are performed in different
6 orders and/or concurrently.

7 Although the description above uses language that is specific to structural
8 features and/or methodological acts, it is to be understood that the invention
9 defined in the appended claims is not limited to the specific features or acts
10 described. Rather, the specific features and acts are disclosed as exemplary forms
11 of implementing the invention.